

Common Knowledge in Interaction Structures*

Krzysztof R. Apt Andreas Witzel
Jonathan A. Zvesper

ILLC, Universiteit van Amsterdam, Netherlands
and CWI, Amsterdam, Netherlands

July 2, 2009

Abstract

We consider two simple variants of a framework for reasoning about knowledge amongst communicating groups of players. Our goal is to clarify the resulting epistemic issues. In particular, we investigate what is the impact of common knowledge of the underlying hypergraph connecting the players, and under what conditions common knowledge distributes over disjunction. We also obtain two versions of the classic result that common knowledge cannot be achieved in the absence of a simultaneous event (here a message sent to the whole group).

1 Introduction

We introduce a framework for reasoning about communication amongst groups of players. We assume that each player is a member of a certain number of groups, and that he is able to broadcast synchronously information to each of those groups. Thus there is what we call an *interaction structure*, a hypergraph of the players, that determines the communication protocol. We are interested in studying what players can learn in certain restricted communication settings, what impact common knowledge of the underlying hypergraph can have, and in properties of the resulting knowledge that can simplify reasoning about it.

For example, consider Figure 1. If player i knows that he is in interaction structure (a), and he learns a fact from player j that initially only player n knew, then i can deduce that both l and k also must have learned that fact. In interaction structure (b), he can only deduce that either of them has learned it, but not which one. If i does not know the interaction structure, he cannot draw such conclusions, since player n might as well have communicated with player j directly. One particular focus of our discussion concerns conditions

*To appear in Proceedings of TARK 2009




Figure 1: Two interaction structures. Hyperarcs are shown in gray.

under which knowledge of a disjunction does allow us to deduce knowledge of one particular disjunct, thus simplifying reasoning in such situations. Another focus is to analyze the conditions for attaining common knowledge.

In the following Section 2, we first set up a more restricted framework where players can only communicate those facts that they initially know, and we examine this framework in detail in Section 3. In Section 4 we then lift this restriction and examine how the properties of knowledge are affected when players are allowed to send information which they learned from other players. In Section 5 we discuss related work, in particular two closely related frameworks from the literature, and draw some conclusions. We look at some possible extensions in Section 6.

2 Preliminaries

We assume the following setup to be common knowledge among the players. There is a set of players N . Each player $i \in N$ has a private set At_i of facts (atomic propositions), of which only player i initially knows whether they are true. The truth values of these facts are represented by a **valuation**, which can be written as a set $V \subseteq At$ containing those facts that are true, where $At = \bigcup_{i \in N} At_i$. By V_i , we denote $V \cap At_i$, the restriction of V to i 's facts.

Throughout, we assume communication to be **truthful** in the sense that it only contains information the sender knows to be true.

An **interaction structure** for players N is a tuple $(H, (At_i)_{i \in N})$, where H is a **hypergraph on** N , i.e., a set of non-empty subsets of N , called **hyperarcs**, and the At_i are pairwise disjoint sets.

In the present section we place two restrictions, that are related. Firstly, we use unordered *sets* of messages, i.e. without any temporal structure, since it only matters whether a given message has been broadcast or not, and not *when* it was broadcast. Secondly, we only allow messages of the form (i, A, p) with $i \in A \in H$ and $p \in At_i$. That is, players only broadcast basic facts that ‘belong’ to them. In Section 4 we partially lift these restrictions, allowing more general forms of broadcast. This in turn means introducing some temporal ordering since if the message (i, A, p) occurs, with $p \notin At_i$, then everybody in A knows that *before* that broadcast there was another broadcast of the form (\cdot, B, p) with

$i \in B$, since otherwise i could not have known p .

Given these restrictions, we consider two different situations: one in which the underlying hypergraph is commonly known amongst the players; and one in which it is not, in the sense that a player knows only the hyperarcs to which he belongs.

In each case an interaction structure defines a communication protocol: each player i can at any point broadcast any true fact $p \in At_i$ to any hyperarc $A \in H$ with $i \in A$. Thus a **message** is a tuple (i, A, p) with $i \in A$ and $p \in At_i$; (i, A, p) is the message in which i communicates among the group A his fact p . **H -compliant messages** are those in which $A \in H$. If the players consider only H -compliant messages possible, then they know the underlying hypergraph H . So if the model allows only H -compliant messages, the underlying hypergraph H is common knowledge among the players; if it uses all messages, H is unknown.

We next define our model formally in order to reason about the knowledge of the players and how it changes as messages are broadcast. This is roughly along the lines of *history based models* (see, e.g., Pacuit and Parikh (2007); Fagin et al. (1995)). We start by defining a **state**, which we might also have called ‘possible world’, (V, M) to consist of a valuation $V \subseteq At$ and a set M of messages (\cdot, \cdot, p) such that $p \in V$. An **H -compliant** state is one where M only contains H -compliant messages.

A **word** over a set $A \subseteq N$ is a finite sequence $w = i_1 \dots i_k$ where each $i_l \in A$. By A^* we denote the set of all words over A , and we write $Set(w)$ for the *set* of players occurring in w .

Now given a set of messages M and a word w , we introduce the following notation:

$$M_w := \{(\cdot, A, \cdot) \in M \mid Set(w) \subseteq A\}$$

$$Facts(M) := \{p \mid (\cdot, \cdot, p) \in M\}.$$

So M_i (respectively, M_w) is the subset of the set of messages M that player i received (respectively, that were broadcast to all the players in w ; note that the order in w does not matter), and $Facts(M)$ is the set of facts that were communicated in the messages in M . In particular, $Facts(M_i)$ is the set of facts that were communicated in the messages in M that player i received. Note that (V, M) is a state if $Facts(M) \subseteq V$. Further, we define all set operations to act component-wise on states, e.g. $(V, M) \subseteq (V', M')$ iff $V \subseteq V'$ and $M \subseteq M'$.

In order to represent the knowledge of the players we define an **indistinguishability relation** between states: $(V, M) \sim_i (V', M')$ iff $(V_i, M_i) = (V'_i, M'_i)$.

In the semantics we present below, a player i is said to ‘know’ a fact just if that fact is true in every state that is indistinguishable for i from the actual state. Of particular interest to us is the knowledge of *groups* $G \subseteq N$ (always assumed to be non-empty). Specifically we consider the so-called ‘common knowledge’ among a group (cf. (Fagin et al., 1995, p. 23)). These are facts that everybody in the group knows, they all know that they know, etc. To define this formally we extend the individual indistinguishability relation to groups: for $G \subseteq N$ the

relation \sim_G is the transitive closure of $\bigcup_{i \in G} \sim_i$.

We are interested in properties definable by the following *epistemic language* \mathcal{L} :

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid C_G\varphi,$$

where the atoms p denote the facts in At , \neg , \wedge and \vee are the standard connectives; and C_G is a knowledge operator, with $C_G\varphi$ meaning φ is common knowledge among G . We write K_i for $C_{\{i\}}$; $K_i\varphi$ can be read ‘ i knows that φ ’. The **positive language** \mathcal{L}^+ is the sublanguage of \mathcal{L} in which negation (\neg) does not occur.

The semantics for \mathcal{L} is as follows:

$$\begin{aligned} (V, M) \models_H p & \quad \text{iff } p \in V, \\ (V, M) \models_H \neg\varphi & \quad \text{iff } (V, M) \not\models_H \varphi, \\ (V, M) \models_H \varphi \vee \psi & \quad \text{iff } (V, M) \models_H \varphi \text{ or } (V, M) \models_H \psi, \\ (V, M) \models_H \varphi \wedge \psi & \quad \text{iff } (V, M) \models_H \varphi \text{ and } (V, M) \models_H \psi, \\ (V, M) \models_H C_G\varphi & \quad \text{iff } (V', M') \models_H \varphi \\ & \quad \text{for each } H\text{-compliant } (V', M') \\ & \quad \text{with } (V, M) \sim_G (V', M'). \end{aligned}$$

By allowing only H -compliant states in the last clause of the semantics, the underlying hypergraph H is assumed to be *common knowledge*. Assuming that the hypergraph H is *unknown* turns out to be equivalent to the case where it is common knowledge that the hypergraph H is complete, i.e., $H = \mathcal{P}(N) - \emptyset$. This might seem counter-intuitive, but it reflects the fact that if the hypergraph is unknown then every player must consider it possible that every set $A \subseteq N$ might be a hyperarc in H . To denote the corresponding semantics, we use \models as abbreviation for \models_H with H being the complete hypergraph.

For a word $w = i_1 \dots i_k$, we write K_w to abbreviate $K_{i_1} K_{i_2} \dots K_{i_k}$, and write \sim_w to denote the concatenation $\sim_{i_1} \circ \dots \circ \sim_{i_k}$.

Notice that $(V, M) \sim_G (V', M')$ iff there is $w \in G^*$ with $(V, M) \sim_w (V', M')$. So an equivalent way of specifying the semantics for C_G with non-singleton G is as follows:

$$(V, M) \models C_G\varphi \text{ iff } (V, M) \models K_w\varphi \text{ for all } w \in G^*. \quad (\star)$$

We now study the consequences of two choices in the analysis of players’ knowledge:

- The type of messages; we assumed already that players send only atomic information, but there still remains a choice whether, as assumed above, players only send information they know initially, or can send information that they have learned from other players. The former scenario is explored in Section 3, the latter in Section 4.
- The issue whether the underlying hypergraph is commonly known among the players. We consider this distinction in both of the following sections.

We shall see that both choices have bearing on players’ knowledge.

3 Telling

In this section we study the case under the assumption mentioned above, that players' messages refer only to the facts they know initially. So players can send only information they know at the outset. We call this contingency 'telling'.

For the relevance of common knowledge of H , consider the following example.

Example 3.1. For players $G = \{i, j, k\}$, $H = \{G\}$, $p \in At_k$ and $(V, M) = (\emptyset, \emptyset)$, we have

$$(V, M) \models_H K_i \neg K_j p.$$

Indeed, the only hyperarc in H through which player j could learn anything from k is the one which also contains player i . So there is no way for k to tell j anything 'secretly'. Hence, with $M_i = \emptyset$, i also knows that $M_{jk} = \emptyset$. That is, in all states which i considers possible at (V, M) , k has not told j that p , therefore in all these states j does not know p .

On the other hand, we have $(V, M) \not\models K_i \neg K_j p$, since $(V, M) \sim_i (\{p\}, \{(k, \{j, k\}, p)\})$.

So for some formulas, common knowledge of H matters. Note also that in this example, we even have

$$(V, M) \models_H C_G \neg K_j p,$$

which shows that common knowledge can be attained without any communication taking place. \square

However, common knowledge of formulas from the positive language \mathcal{L}^+ can only be attained through messages received by the whole group, and for these formulas, common knowledge of H does not matter. In order to establish this, we first show the following Lemma 3.2, which intuitively says that if a positive formula is true in some state, then it remains true in any state where more facts are true or more communication has taken place. Remember that \models corresponds to \models_H with H being the complete hypergraph, so the following carries over to general states and \models .

Lemma 3.2. *For any $\varphi \in \mathcal{L}^+$ and H -compliant states (V, M) and (V', M') with $(V, M) \subseteq (V', M')$,*

$$\text{if } (V, M) \models_H \varphi, \text{ then } (V', M') \models_H \varphi.$$

Proof. We proceed by structural induction on φ . The only not completely obvious case is when $\varphi = C_G \psi$ with $\psi \in \mathcal{L}^+$. We show the claim for $G = \{i\}$; the non-singleton case then follows by induction and (\star) . Take an H -compliant state (V'', M'') such that $(V', M') \sim_i (V'', M'')$. Let

$$(V''', M''') := (V_i \cup \bigcup_{j \neq i} V_j'', M_i).$$

We have $Facts(M_i) \subseteq Facts(M) \subseteq V$, since (V, M) is a state. Also, $M_i \subseteq M'_i = M''_i \subseteq M''$, so $Facts(M_i) \subseteq Facts(M'') \subseteq V''$, since (V'', M'') is a state.

Hence, $Facts(M''') = Facts(M_i) \subseteq V \cap V'' = \bigcup_{i \in N} (V_i \cap V_i'') \subseteq V'''$. This shows that (V''', M''') is a state. Moreover, $(V, M) \sim_i (V''', M''')$. Assume now $(V, M) \models_H K_i \psi$. Then we obtain $(V''', M''') \models_H \psi$. Further, we have $(V''', M''') \subseteq (V'', M'')$ since $V_i \subseteq V_i' = V_i''$ and $M_i \subseteq M_i' = M_i''$ due to $(V', M') \sim_i (V'', M'')$. Thus, by induction hypothesis we obtain $(V'', M'') \models_H \psi$. \square

Theorem 3.3. *For any H -compliant state (V, M) and $\varphi \in \mathcal{L}^+$,*

$$(V, M) \models \varphi \text{ iff } (V, M) \models_H \varphi.$$

Proof. We proceed by structural induction. The only non-trivial step is when $\varphi = C_G \psi$ with $\psi \in \mathcal{L}^+$.

(\Rightarrow) By induction hypothesis, $(V, M) \models C_G \psi$ implies $(V, M) \models_H C_G \psi$, since each H -compliant state is also a state.

(\Leftarrow) Assume to the contrary that $(V, M) \not\models C_G \psi$. So there is a state (V', M') with $(V, M) \sim_G (V', M')$ and $(V', M') \not\models \psi$. Now let

$$M' \upharpoonright_H := \{(\cdot, A, \cdot) \in M' \mid A \in H\}.$$

So $M' \upharpoonright_H$ consists of all H -compliant messages of M' . Now note that $(V', M' \upharpoonright_H) \subseteq (V', M')$, so from $(V', M') \not\models \psi$ we obtain that $(V', M' \upharpoonright_H) \not\models \psi$ using Lemma 3.2 (which, as noted, also holds for general states and \models). Since $(V', M' \upharpoonright_H)$ is H -compliant, the induction hypothesis yields $(V', M' \upharpoonright_H) \not\models_H \psi$. Moreover, we also have $(V, M) \sim_G (V', M' \upharpoonright_H)$, since (V, M) is H -compliant and $(V, M) \sim_G (V', M')$. Thus, $(V, M) \not\models_H C_G \psi$. \square

In the remainder of this section, we are concerned with formulas from \mathcal{L}^+ , so in view of the above results we restrict attention to \models .

We now establish that C_G distributes over disjunctions of positive formulas, starting with singleton G .

Lemma 3.4. *For any $\varphi_1, \varphi_2 \in \mathcal{L}^+$, $i \in N$, and state (V, M) ,*

$$(V, M) \models K_i(\varphi_1 \vee \varphi_2) \text{ iff } (V, M) \models K_i\varphi_1 \vee K_i\varphi_2.$$

Proof. To deal with the (\Rightarrow) implication assume that $(V, M) \not\models K_i\varphi_1 \vee K_i\varphi_2$. Then $(V, M) \not\models K_i\varphi_1$ and $(V, M) \not\models K_i\varphi_2$, i.e., there are (V', M') and (V'', M'') such that

$$\begin{aligned} (V, M) \sim_i (V', M') \text{ and } (V', M') \not\models \varphi_1, \text{ as well as} \\ (V, M) \sim_i (V'', M'') \text{ and } (V'', M'') \not\models \varphi_2. \end{aligned}$$

Let now

$$(V''', M''') := (V_i \cup \bigcup_{j \neq i} (V_j' \cap V_j''), M_i).$$

Then $Facts(M) \subseteq V$, $Facts(M') \subseteq V'$, $Facts(M'') \subseteq V''$, since (V, M) , (V', M') and (V'', M'') are states. Moreover, $M_i = M'_i$ and $M_i = M''_i$. So,

$$\begin{aligned} Facts(M''') &\subseteq Facts(M) \cap Facts(M') \cap Facts(M'') \\ &\subseteq V \cap V' \cap V'' \\ &= \bigcup_{i \in N} (V_i \cap V'_i \cap V''_i) \\ &\subseteq V'''. \end{aligned}$$

This shows that (V''', M''') is a state, and since $M''' = M_i \subseteq M$ it is H -compliant.

Now since $V_i = V'_i = V''_i$ and $M_i = M'_i = M''_i$, we have $(V''', M''') \subseteq (V', M')$ and $(V''', M''') \subseteq (V'', M'')$. By Lemma 3.2, we obtain $(V''', M''') \not\models \varphi_1$ and $(V''', M''') \not\models \varphi_2$, thus $(V''', M''') \not\models \varphi_1 \vee \varphi_2$. Furthermore $(V, M) \sim_i (V''', M''')$, so $(V, M) \not\models K_i(\varphi_1 \vee \varphi_2)$.

Further, the (\Leftarrow) implication immediately holds by the semantics. \square

Theorem 3.5. *For any $\varphi_1, \varphi_2 \in \mathcal{L}^+$, state (V, M) , and $G \subseteq N$,*

$$(V, M) \models C_G(\varphi_1 \vee \varphi_2) \text{ iff } (V, M) \models C_G\varphi_1 \vee C_G\varphi_2.$$

Proof. The claim follows directly from Lemma 3.4 and (\star) . \square

To see that this result does not hold if we allow negation, consider three players $i, j, k \in N$, $p \in V_k$, $(V, M) = (\emptyset, \emptyset)$, and $\varphi = K_i(K_j p \vee \neg K_j p)$. Then $(V, M) \models \varphi$, since the used disjunction is a tautology, but there is no way for i to know which disjunct is true.

Even with non-tautological disjunctions, the result does not hold.

Example 3.6. With $p \in V_k$ and

$$\begin{aligned} (V, M) &= (\{p\}, \{(k, \{i, k\}, p)\}) \\ \varphi &= K_i(K_j p \vee \neg(K_j p \vee K_j \neg p)), \end{aligned}$$

we have $(V, M) \models \varphi$, but again, i knows neither disjunct in (V, M) . Intuitively, having privately learned that p is true, i knows that j either also learned it, or that j doesn't know whether p is true, but i does not know which of these two statements is true. \square

Another observation is that mutual knowledge of any *fact* $p \in At$ can only be obtained through a corresponding message, and is thus inseparably tied to common knowledge.

Lemma 3.7. *For any $w \in N^*$ with $|Set(w)| \geq 2$, $p \in At$, and state (V, M) , the following are equivalent:*

- (i) $(V, M) \models K_w p$,
- (ii) there is $(\cdot, \cdot, p) \in M_w$,

(iii) $(V, M) \models C_G p$ with $G = \text{Set}(w)$.

Proof. (i) \Rightarrow (ii): Assume that (ii) does not hold. Let $V' := V \setminus \{p\}$. Then (V', M_w) is a state and $(V', M_w) \not\models p$. Now let $i \in N$ be such that $p \in \text{At}_i$. Since $|\text{Set}(w)| \geq 2$, there is $j \in \text{Set}(w)$ with $p \notin \text{At}_j$. By construction, $(V', M_w) \not\models p$ and $(V, M) \sim_j (V', M_w)$. Since $j \in \text{Set}(w)$ and both $(V, M) \sim_k (V, M)$ and $(V', M_w) \sim_k (V', M_w)$ for all $k \in N$, we obtain $(V, M) \sim_w (V', M_w)$ and thus $(V, M) \not\models K_w p$.

(ii) \Rightarrow (iii): Suppose that $G = \text{Set}(w)$ and take $m \in M_w$. Consider (V', M') such that $(V, M) \sim_G (V', M')$. This means that for a sequence i_1, \dots, i_k of players from G and some states $(V^1, M^1), \dots, (V^k, M^k)$ we have $(V, M) \sim_{i_1} (V^1, M^1) \sim_{i_2} \dots \sim_{i_k} (V^k, M^k)$, where $(V', M') = (V^k, M^k)$. But $i_1 \in G$, so $m \in M_{i_1}$, and consequently $m \in M_{i_1}^1$. Also $i_2 \in G$, so $m \in M_{i_2}^1$, and consequently $m \in M_{i_2}^2$. Continuing this way we conclude that $m \in M_{i_k}^k$, that is $m \in M'_{i_k}$.

Hence, $m \in M'$. This shows that $M_w \subseteq M'$. So $\text{Facts}(M_w) \subseteq V'$, since (V', M') is a state. But by the assumption $p \in \text{Facts}(M_w)$, so $(V', M') \models p$. This proves $(V, M) \models C_G p$.

(iii) \Rightarrow (i): By (\star) . \square

We can extend this connection between mutual and common knowledge to arbitrary positive formulas.

Theorem 3.8. *For any $G \subseteq N$, $\varphi \in \mathcal{L}^+$, and state (V, M) ,*

$$(V, M) \models C_G \varphi \text{ iff } (V, M) \models K_w \varphi \text{ for some } w \in G^* \text{ with } \text{Set}(w) = G.$$

Proof. The direction (\Rightarrow) is by (\star) .

For (\Leftarrow) , we proceed by structural induction. The base case is obtained from Lemma 3.7. The induction step for disjunction follows by Theorem 3.5, and for conjunction it follows directly by definition of the semantics. For $\varphi = K_i \psi$, the assumption $(V, M) \models K_w K_i \psi$ yields, by induction hypothesis, that $(V, M) \models C_{G'} \psi$ for $G' = \text{Set}(w) \cup \{i\} = G \cup \{i\}$, which by definition of the semantics implies that $(V, M) \models C_G K_i \psi$. \square

Note that this result provides for positive formulas a simplified characterization of the common knowledge operator, as compared with (\star) .

Finally, we establish a result intuitively saying that a group's common knowledge of a positive formula can only be achieved when some message (or messages) has been broadcast to at least all members of this group. So common knowledge of a positive formula cannot be achieved among a group by means of more limited communications, for example point-to-point messages. Given a formula φ we denote by $\text{Facts}(\varphi)$ the set of facts that occur in it.

Theorem 3.9. *For any $G \subseteq N$ with $|G| \geq 2$, $\varphi \in \mathcal{L}^+$, and state (V, M) ,*

if $(V, M) \models C_G \varphi$, then there is $(\cdot, A, p) \in M$ with $G \subseteq A$ and $p \in \text{Facts}(\varphi)$.

Proof. By Theorem 3.5 and the definition of semantics, we can transform $C_G\varphi$ into an equivalent formula consisting only of disjunctions and conjunctions over formulas of the form $C_GC_{G_1}\dots C_{G_l}p$ with $G_1, \dots, G_l \subseteq N$ and $p \in \text{Facts}(\varphi)$. Since $(V, M) \models C_G\varphi$ there is at least one of these formulas for which $(V, M) \models C_GC_{G_1}\dots C_{G_l}p$.

Take now w such that $\text{Set}(w) = G$. By (\star) we obtain $(V, M) \models K_w C_GC_{G_1}\dots C_{G_l}p$, so by the definition of semantics $(V, M) \models K_w p$. The claim now follows by Lemma 3.7. \square

4 Forwarding

We now consider a more complex situation in which players are allowed to send facts that they learned from other players. We call this contingency ‘forwarding’. It is achieved by relaxing in the definition of a message (i, A, p) the assumption $p \in \text{At}_i$ to $p \in \text{At}$. We still insist that a player can send a message only to a group to which he belongs, that is, that $i \in A$ holds.

We also assume that only information *known to be true* is sent, so we now need to examine how a player learned the information he is sending. This brings us to consider the following relation on the set of messages:

$$(j, B, p) \rightsquigarrow (i, A, p) \text{ iff } p \notin \text{At}_i \text{ and } i \in B.$$

Intuitively, $(j, B, p) \rightsquigarrow (i, A, p)$ means that the fact p is initially not known to player i and that he has learned it from a message sent by player j to a group to which i belongs. So $(j, B, p) \rightsquigarrow (i, A, p)$ means that (j, B, p) is a possible (partial) *explanation* of (i, A, p) .

By a **state** we now mean a pair (V, M) such that for each message $(i, A, p) \in M$ a sequence of messages $(j_1, B_1, p), \dots, (j_k, B_k, p)$ exists (i.e., each of these messages about p is in M) such that

- these messages form an explanatory chain: for $l \in \{1, \dots, k-1\}$ we have $(j_l, B_l, p) \rightsquigarrow (j_{l+1}, B_{l+1}, p)$;
- they are not circular: players j_1, \dots, j_k are all different;
- p is initially known to player j_1 : $p \in V_{j_1}$; and
- the fact p reaches player i : $(j_k, B_k, p) = (i, A, p)$.

We call such a sequence of messages an **explanation** for (i, A, p) in (V, M) . So a pair (V, M) is a state if for each of its messages it has an explanation.

Note that given a state, its messages contain only true facts. That is, if (V, M) is a state, then $\text{Facts}(M) \subseteq V$. Moreover, if $(i, A, p) \in M$, then player i knows that p is true, i.e. $(V, M) \models K_i p$. (A more general statement is established in Lemma 4.5.) Note also that when each message in M is of the form (i, \cdot, p) , where $p \in V_i$, then (V, M) is a state, since each message then forms its own explanation. So states considered in this section generalize the states considered in the previous section.

Each state can be alternatively viewed as a partial ordering \rightsquigarrow^* (where \rightsquigarrow^* is the reflexive, transitive closure of the \rightsquigarrow relation) on a set of messages such that each message has an explanation.

The only restriction on the order of the actions comes from the relation \rightsquigarrow that needs to be respected: a player sends a message that contains information that either he initially knows to be true (the message is (i, A, p) where $p \in V_i$) or he has learned (the message is (i, A, p) and some earlier message is of the form (j, B, p) , where $i \in A$). So the computation begins by some players who send information they know is true.

We now consider the semantics introduced in Section 2 in this extended setting. It is important to realize that these two semantics differ in the sense that for a state and a formula $\varphi \in \mathcal{L}$ it can happen that $(V, M) \models_H \varphi$ holds in the sense of Section 2 but not in the sense considered now.

Example 4.1. Let $N = \{i, j, k\}$, $H = \{\{i, j\}, \{j, k\}\}$, $V = \{p\}$, where $p \in At_i$, and $M = \{(i, \{i, j\}, p)\}$. Then we have

$$(V, M) \models_H K_i \neg K_k p$$

in the sense of Section 2. The intuitive reason is that the fact p ‘belongs’ to i , so it cannot be used in any message sent by j , and this information is known to i . However, in the present setting p can be used in a message sent by j and we have

$$(V, M) \not\models_H K_i \neg K_k p.$$

Indeed, consider (V', M') with $V' = \{p\}$ and $M' = \{(i, \{i, j\}, p), (j, \{j, k\}, p)\}$. Then $(V, M) \sim_i (V', M')$ and $(V', M') \models K_k p$. \square

In general, only non-epistemic formulas have the same meaning w.r.t. both semantics.

We now show that some, though not all, properties established in the previous section also hold in this new setting. In particular, as in Example 3.1, we have $(\emptyset, \emptyset) \models_H C_G \neg K_j p$, so also now common knowledge can exist without any communication taking place. However, as we shall see, Theorem 3.3 does not hold in general any more. So in the following we usually consider \models_H , which, as mentioned earlier, includes \models as a special case with H being the complete hypergraph.

Recall that for a formula φ we denoted the set of facts that occur in it by $Facts(\varphi)$.

Lemma 4.2. *For any $\varphi \in \mathcal{L}^+$ and H -compliant state (V, M) if $(V, M) \models_H \varphi$, then $Facts(\varphi) \cap V \neq \emptyset$.*

Proof. We proceed by structural induction on φ . The only not completely obvious case is when $\varphi = C_G \psi$. But $(V, M) \models_H C_G \psi$ implies $(V, M) \models_H \psi$, so in this case the induction hypothesis readily applies, as well. \square

The following result is then a counterpart of Theorem 3.9.

Theorem 4.3. *For any $G \subseteq N$ with $|G| \geq 2$, $\varphi \in \mathcal{L}^+$, and H -compliant state (V, M) ,*

if $(V, M) \models_H C_G \varphi$, then there is $(\cdot, A, p) \in M$ with $G \subseteq A$ and $p \in Facts(\varphi)$.

Proof. Note that the conclusion of the implication can be written in a more succinct way as $Facts(\varphi) \cap Facts(\bigcap_{i \in G} M_i) \neq \emptyset$. Suppose that $(V, M) \models_H C_G \varphi$ and $Facts(\varphi) \cap Facts(\bigcap_{i \in G} M_i) = \emptyset$. Call a message a *p-message* if it is of the form (\cdot, \cdot, p) . Abbreviate $\bigcup_{i \in G} M_i$ to M_G (Note that this is different from M_w , where w is a word, which corresponds to an intersection.) Three cases arise.

Case 1. For all $p \in Facts(\varphi)$ there is no p -message in M_G and $Facts(\varphi) \cap V = \emptyset$. Then by Lemma 4.2 $(V, M) \not\models_H \varphi$.

Case 2. For all $p \in Facts(\varphi)$ there is no p -message in M_G and $Facts(\varphi) \cap V \neq \emptyset$.

Take some $p \in Facts(\varphi) \cap V$. Since $|G| \geq 2$, there is $i \in G$ such that $p \notin V_i$. Remove from V the fact p and from M all p -messages. Denote the outcome by (V', M') . By construction (V', M') is an H -compliant state and by the assumption there is no p -message in M_i , so $(V, M) \sim_i (V', M')$.

Case 3. For some $p \in Facts(\varphi)$ there is a p -message in M_G .

Given a set of messages $O \subseteq M$, we denote by $top(O)$ the set of p -messages $m \in O$, where $p \in Facts(\varphi)$, such that for no $m' \in O$ we have $m \neq m'$ and $m \rightsquigarrow^* m'$. Further, we define (cl stands for the closure)

$$cl(m) := \{m' \in M \mid m \rightsquigarrow^* m'\}.$$

We assumed that the set of p -messages in M_G , where $p \in Facts(\varphi)$, is non-empty, so the set $top(M_G)$ is non-empty and hence for some $i_1 \in G$ the set $top(M_G) \cap M_{i_1}$ is non-empty. Choose some $m \in top(M_G) \cap M_{i_1}$. Let $M' := M \setminus cl(m)$ and $V' := V$. By the assumption $Facts(\varphi) \cap Facts(\bigcap_{i \in G} M_i) = \emptyset$, so there is some $i \in G$ such that $i \notin A$, where m is of the form (\cdot, A, p) for some $p \in Facts(\varphi)$. By the construction (V', M') is an H -compliant state and $(V, M) \sim_i (V', M')$.

We now repeat the above case analysis with (V', M') instead of (V, M) . Iterating this way we eventually end up in Case 1 since in Cases 2 and 3 always some fact or message is removed. This way we obtain a word $w \in G^*$ and an H -compliant state (V'', M'') such that $(V, M) \sim_w (V'', M'')$ and $(V'', M'') \not\models_H \varphi$. By (\star) this contradicts the assumption that $(V, M) \models_H C_G \varphi$. \square

Corollary 4.4. *For any $G \subseteq N$ with $|G| \geq 2$, $p \in At$, and H -compliant state (V, M) ,*

$$(V, M) \models_H C_G p \text{ iff there is } (\cdot, A, p) \in M \text{ with } G \subseteq A.$$

Proof. (\Rightarrow) is a direct consequence of Theorem 4.3.

(\Leftarrow) The proof is analogous to the one of the implication $(ii) \Rightarrow (iii)$ of Lemma 3.7 and is omitted. \square

Here is the counterpart of the above result for the case of one player. It states that in any state a player knows a fact iff either he knows it at the outset or he has learned it through a message he received.

Lemma 4.5. *For any $i \in N$, $p \in At$, and H -compliant state (V, M) ,*

$$(V, M) \models_H K_i p \text{ iff } p \in V_i \cup Facts(M_i).$$

Proof. (\Rightarrow) Suppose that $(V, M) \models_H K_i p$ and $p \notin V_i \cup \text{Facts}(M_i)$. Remove from V the fact p and from M all messages of the form (\cdot, \cdot, p) . Denote the outcome by (V', M') . By construction (V', M') is an H -compliant state and by the assumption $(V, M) \sim_i (V', M')$. So $(V', M') \models_H p$, which is a contradiction.

(\Leftarrow) Consider an H -compliant state (V', M') such that $(V, M) \sim_i (V', M')$. Then $V_i = V'_i$ and $M_i = M'_i$. So $V_i \subseteq V'$ and $\text{Facts}(M_i) \subseteq \text{Facts}(M') \subseteq V'$, where the final inclusion follows by the fact that (V', M') is a state. So $p \in V'$ and consequently $(V', M') \models_H p$, as desired. \square

Corollary 4.6. *For any $i \in N$, $p \in \text{At}$, and H -compliant state (V, M) ,*

$$(V, M) \models_H K_i p \text{ iff } (V, M) \models K_i p.$$

Proof. This follows from Lemma 4.5 and the fact that \models is a special case of \models_H . \square

For further analysis we need an auxiliary concept. Suppose that $(V, M) \subseteq (V', M')$, where (V', M') is a state. In general, (V, M) does not need to be a state but we can complete it to a state $L(V, M)$ such that $L(V, M) \subseteq (V', M')$. Indeed, it suffices for each message m in M to add to M messages forming an explanation of m in M' and then add to V the facts used in these added messages. More precisely, let M'' be a smallest set such that $M \subseteq M'' \subseteq M'$ and $(V \cup \text{Facts}(M''), M'')$ is a state. In general, this does not define a unique state, since each message in M' can have multiple explanations. However, the states are finite, so we can always choose $(V \cup \text{Facts}(M''), M'')$ in a unique way, for example, by associating with each state a unique natural number.

From now on we assume that given an inclusion $(V, M) \subseteq (V', M')$ the state $L(V, M)$ is uniquely defined. Note that if (V', M') is H -compliant, then so are (V, M) and $L(V, M)$.

The following observation will be useful.

Fact 4.7. *For any $i \in N$ and H -compliant states $(V, M), (V', M')$ with $(V_i, M_i) \subseteq (V'_i, M'_i)$,*

$$(V, M) \sim_i L(V_i, M_i).$$

Proof. All messages in (V_i, M_i) involve player i , so the H -compliant state $(V'', M'') = L(V_i, M_i)$ is realized by adding to (V_i, M_i) only some messages that do not involve player i and some facts from outside of At_i . Consequently $M_i = M''_i$ and $V_i = V''_i$, that is $(V, M) \sim_i (V'', M'')$. \square

Next, the following property of the semantics will be needed.

Lemma 4.8. *For any H -compliant state (V, M) , $G \subseteq N$, and facts p_1, \dots, p_k ,*

$$(V, M) \models_H C_G(\bigvee_{j=1}^k p_j) \text{ iff } (V, M) \models_H \bigvee_{j=1}^k C_G p_j.$$



Figure 2: Knowledge of H matters even for positive formulas when forwarding is allowed.

Proof. To deal with (\Rightarrow) we consider two cases.

Case 1. $|G| = 1$, say $G = \{i\}$.

Suppose that $(V, M) \not\models_H \bigvee_{j=1}^k K_i p_j$. Then for $j \in \{1, \dots, k\}$ we have $(V, M) \not\models_H K_i p_j$ and thus by Lemma 4.5 $p_j \notin V_i \cup \text{Facts}(M_i)$. Let

$$(V'', M'') := L(V_i, M_i)$$

be the H -compliant state defined w.r.t. the inclusion $(V_i, M_i) \subseteq (V, M)$. This state is realized by adding to M_i some messages from M and to V_i some facts from $\text{Facts}(M_i)$. So $V'' \subseteq V_i \cup \text{Facts}(M_i)$ and consequently for $j \in \{1, \dots, k\}$ we have $p_j \notin V''$. Hence,

$$(V'', M'') \not\models_H \bigvee_{j=1}^k p_j.$$

Moreover, by Fact 4.7 we have $(V, M) \sim_i (V'', M'')$, so $(V, M) \not\models_H K_i(\bigvee_{j=1}^k p_j)$.

Case 2. $|G| \geq 2$.

By Lemma 4.2 for some $j \in \{1, \dots, k\}$ there is $(\cdot, A, p_j) \in M$ with $G \subseteq A$. So, by Corollary 4.4, $(V, M) \models_H C_G p_j$, and thus $(V, M) \models_H \bigvee_{j=1}^k C_G p_j$.

The (\Leftarrow) implication holds directly by the definition of the semantics. \square

We can now resume our comparison with the results of the previous section. To start with, the following result is a counterpart of Lemma 3.2.

Lemma 4.9. *For any $\varphi \in \mathcal{L}^+$ and H -compliant states (V, M) and (V', M') with $(V', M') \subseteq (V, M)$,*

$$\text{if } (V', M') \models_H \varphi, \text{ then } (V, M) \models_H \varphi.$$

Proof. By structural induction on φ . \square

In Section 3 we used this result to establish Theorem 3.3. However, in the current setting the counterpart of Theorem 3.3 does not hold.

Example 4.10. Consider players $N = \{i, j, k, l\}$ and a graph H with the edges $\{l, k\}, \{k, j\}, \{j, i\}$, see Figure 2. Suppose that $V = \{p\}$, where $p \in \text{At}_l$, and $M = \{(l, \{l, k\}, p), (k, \{k, j\}, p), (j, \{j, i\}, p)\}$. Then

$$(V, M) \not\models K_i K_k p,$$

since player i does not know through which source player j learned p . However,

$$(V, M) \models_H K_i K_k p,$$

since when the underlying graph is commonly known, player i knows that player j learned p from player k . \square

Still, a limited counterpart of Theorem 3.3 does hold. Let \mathcal{L}_K^+ be the sublanguage of \mathcal{L}^+ in which the knowledge operators C_G are not allowed to be nested. So if $C_G\varphi \in \mathcal{L}_K^+$, then φ is a propositional formula that does not use negation.

Theorem 4.11. *For any H -compliant state (V, M) and $\varphi \in \mathcal{L}_K^+$,*

$$(V, M) \models \varphi \text{ iff } (V, M) \models_H \varphi.$$

Proof. We proceed by structural induction on φ . The only non-trivial case is when $\varphi = C_G\psi$ for some $i \in N$ and ψ is a propositional formula that does not use negation.

Let $\bigwedge_{j=1}^k \bigvee_{l=1}^{m_j} p_{j,l}$ be the conjunctive normal form of ψ . So each $p_{j,l}$ is a fact. By Lemma 4.8 and the definition of semantics we have both

$$(V, M) \models C_G\psi \text{ iff } (V, M) \models \bigwedge_{j=1}^k \bigvee_{l=1}^{m_j} C_G p_{j,l}$$

and

$$(V, M) \models_H C_G\psi \text{ iff } (V, M) \models_H \bigwedge_{j=1}^k \bigvee_{l=1}^{m_j} C_G p_{j,l}.$$

But by Corollary 4.6, for all $j \in \{1, \dots, k\}$ and $l \in \{1, \dots, m_j\}$ we have

$$(V, M) \models C_G p_{j,l} \text{ iff } (V, M) \models_H C_G p_{j,l}.$$

This implies the claim for $C_G\psi$. \square

We now analyze to what extent Theorem 3.5 holds in the current setting. We first prove that the C_G operator distributes over disjunctions of formulas from the non-epistemic sublanguage $\mathcal{L}_{\wedge, \vee}$ of \mathcal{L} in which only conjunction and disjunction is allowed.

Theorem 4.12. *For any $\varphi_1, \varphi_2 \in \mathcal{L}_{\wedge, \vee}$, $i \in N$ and H -compliant state (V, M) ,*

$$(V, M) \models_H C_G(\varphi_1 \vee \varphi_2) \text{ iff } (V, M) \models_H C_G\varphi_1 \vee C_G\varphi_2.$$

Proof. Passing by the conjunctive normal forms of φ_1 and φ_2 , the result follows from the definition of the semantics and Lemma 4.8 twice. \square

However, the C_G operator does not distribute over the knowledge operators, so the counterpart of Theorem 3.5 does not hold.

Example 4.13. Consider the set of players $N = \{i, j, k, l, n\}$ and the hypergraph H being the graph with the edges $(n, k), (n, l), (k, j), (l, j), (j, i)$, see Figure 1(b). Take $V = \{p\}$, where $p \in At_n$, and

$$M = \{(n, \{n, k\}, p), (k, \{k, j\}, p), (j, \{j, i\}, p)\}.$$

Then

$$(V, M) \models_H K_i(K_k p \vee K_l p),$$

but neither $(V, M) \models_H K_i K_k p$, nor $(V, M) \models_H K_i K_l p$ holds. Informally, player i knows that either player k or player l knows p but he does not know which one of them knows p . \square

As noticed already after the proof of Theorem 3.5, the K_i operator does not distribute over negation either; the same example applies here.

Finally, reconsider Theorem 3.8. It is straightforward to see that it does not hold in the present setting, even for two players. Indeed, reconsider Example 4.10. We showed there that $(V, M) \models_H K_i K_k p$. However, it is easy to see that $(V, M) \not\models_H C_{\{i,k\}} p$ since $(V, M) \not\models_H K_k K_i p$.

5 Conclusions and related work

In this paper we studied various aspects of common knowledge in two simple frameworks concerned with synchronous communication. It is useful to clarify that our two impossibility results concerning the attainment of common knowledge amongst players (Theorems 3.9 and 4.3) differ from the customary impossibility results.

For example, Halpern and Moses (1990) formalize the epistemic aspects of the celebrated Coordinated Attack Problem that consists in achieving common knowledge (a ‘common plan of action’). They show (in Section 8) that in a distributed system in which communication is not guaranteed, common knowledge is not attainable. When communication is guaranteed, they show the same result when there is no bound on message delivery times. In both situations the proof assumes the existence of clocks and point-to-point communication.

The close correspondence between simultaneous events (in our system a broadcast to the whole group) and common knowledge is pointed out by Fagin et al. (1999). Their model of a distributed system consists of a set of linear ‘runs’ (histories), while we only assume a partial ordering (\rightsquigarrow) between messages broadcast to groups, which are the only possible actions. We have shown that in our framework, common knowledge of a positive formula is indeed inseparably related to group communication, which corresponds to simultaneous events. However, as we have seen, this does not hold of negative formulas, so the relationship is not as obvious as it may seem. The results of Fagin et al. (1999) may be seen to correspond to our Corollary 4.4, though we allow broadcasts instead of just point-to-point communication.

Chandy and Misra (1986) consider the flow of information in distributed systems with asynchronous communication. They study how processes ‘learn’ about states of other processes and how knowledge evolves. The main difference is that with asynchronous communication, hypergraphs are equivalent to mere point-to-point graphs. Without guarantees on the delivery time, and without temporal reasoning, from the knowledge point of view sending an asynchronous

group message has the same effect as sending a separate message to each group member.

Our study concerning the consequences of the assumption whether the underlying hypergraph is commonly known among the players brings our paper somewhat closer to the area of social networks (see, e.g., Jackson (2008)). Within logic, the relevance of epistemic issues in communication networks has been recognized by a number of authors, e.g. van Benthem (2006). However, to our knowledge the only work that addresses these issues is Pacuit and Parikh (2007) and, to some extent, Roelofsen (2005). We now briefly discuss these frameworks and relate them to our own.

Pacuit and Parikh (2007) use a history-based model to study diffusion of information in a communication graph, starting from facts initially known to individual players. Communicative acts are assumed to consist in a player j ‘reading’ an arbitrary propositional formula from another player i , with the precondition that i *knows* that the formula holds. Communicative acts are restricted to a commonly known, static, directed graph, and, unlike in our case, are assumed to go *unnoticed by* i . The paper formalizes what conclusions, beyond the mere factual content of messages, can be drawn using knowledge of the communication graph and, consequently, knowledge of the possible routes along which certain information can have flown.

Roelofsen (2005) uses a model based on Dynamic Epistemic Logic (DEL) to describe how some initial epistemic model evolves in a communication situation. Communication is among subgroups and can contain arbitrary epistemic formulas. Further, communication is assumed to be truthful and is restricted to occur along a hypergraph. However the hypergraph is explicitly encoded in the model, and thus (knowledge of it) is subject to change.

While under certain circumstances history-based modeling and DEL are equivalent (van Benthem et al., 2007), our approach is more in the spirit of Pacuit and Parikh (2007). Indeed, we also study how specific information may have spread. Also, all possible communications are included in the model and suspicions about them are not explicitly formed. Finally, the underlying graph (in our case hypergraph) is static and not included in the model.

On a technical level, our approach differs from Pacuit and Parikh (2007) in that we use sets of messages instead of sequences and, when dealing with forwarding, employ a more general structure than histories by considering messages partially ordered by the relation \rightsquigarrow . On the other hand, our messages are simpler: Pacuit and Parikh (2007) allow disjunctions of facts, while we allow only facts.

What distinguishes our approach on a more conceptual level is that our focus lies on identifying natural conditions that allow us to prove stronger results about knowledge, such as distributivity over disjunctions, or irrelevance of (common) knowledge of the underlying hypergraph.

6 Extensions

We conclude by listing a number of natural extensions of the considered framework that are worthy of further study:

- We could equip the players with theories that their parts of valuations, V_i , have to satisfy. In this extension we would assume that each player i has a propositional theory T_i built from facts in At_i that he adheres to. The theories T_i where $i \in N$ then form a common knowledge among the players. So each player j can assume that player i considers V_i such that V_i is a model of T_i .
- We could consider more complex messages than simple atomic facts, for example propositional formulas, or even epistemic formulas. Also, we could study asynchronous communication, messages from unknown senders or to an unknown group of recipients, and a counterpart of the blind copy feature familiar from e-mails.
- In Section 4 we relaxed the assumption that in a message (i, A, p) it has to be the case that $p \in At_i$, but we did still insist on the *truthfulness* of messages, requiring that $p \in V$. We could further relax this assumption, by insisting only that $p \in At$. This way we would model messages that consist of possibly false (but credible) information. This would lead to a study of beliefs (which can be false) rather than knowledge (which cannot) and common beliefs rather than common knowledge.
- We could consider in this framework belief revision, by assuming that the theory T_i of player i consists of his beliefs, which would then be revised in view of received information. Alternatively, T_i could be the certain knowledge of player i against which received information would be revised.
- We could assume that the players have different knowledge of the underlying hypergraph, by assuming that for all i we have $H \subseteq H_i$, where H is the underlying hypergraph and H_i is its approximation known to i , and that players learn H by exchanging messages. The messages would contain information about which hyperarcs do *not* belong to H .
- Alternatively, we could study a setup in which each player has an indistinguishability relation over hypergraphs. This would allow us to model players' partial knowledge of the underlying hypergraph.

We use the setting of the first item in (Apt et al., 2009) to reason about iterated elimination of strategies in ***strategic games with interaction structures***. These are strategic games in which there is a hypergraph over the set of players (an interaction structure) and the players can communicate about their preferences, initially only known to themselves, so that within each hyperarc players can obtain common knowledge of each other's preferences.

Acknowledgements

We thank Rohit Parikh, Willemien Kets, Aaron Archer, Henry Landau, and three anonymous referees for discussion and helpful suggestions. The second and third authors were supported by a GLoRiClass fellowship funded by the European Commission (Early Stage Research Training Mono-Host Fellowship MEST-CT-2005-020841).

References

- K. R. Apt, A. Witzel, and J. A. Zvesper. Strategy elimination in games with interaction structures. Manuscript, 2009.
- K. M. Chandy and J. Misra. How processes learn. *Distributed Computing*, 1(1): 40–52, Mar. 1986.
- R. Fagin, J. Halpern, M. Vardi, and Y. Moses. *Reasoning about knowledge*. MIT Press, Cambridge, MA, USA, 1995.
- R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. Common knowledge revisited. *Annals of Pure and Applied Logic*, 96(1–3):89–105, 1999.
- J. Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, 1990.
- M. O. Jackson. *Social and Economic Networks*. Princeton University Press, Princeton, 2008.
- E. Pacuit and R. Parikh. Reasoning about communication graphs. In *Interactive Logic*, volume 1 of *Texts in Logic and Games*, pages 135–157, London, 2007. Amsterdam University Press.
- F. Roelofsen. Exploring logical perspectives on distributed information and its dynamics. Master’s thesis, ILLC, University of Amsterdam, 2005.
- J. van Benthem. One is a lonely number: Logic and communication. In *Logic Colloquium 2002*, volume 27 of *Lecture Notes in Logic*, pages 96–129. A K Peters, 2006.
- J. van Benthem, J. Gerbrandy, and E. Pacuit. Merging frameworks for interaction: DEL and ETL. In *Proceedings of the 11th conference on Theoretical aspects of rationality and knowledge*, pages 72–81, Brussels, Belgium, 2007. ACM.